

Kaxiras 2-2

AF
JFW

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Patent Application

Applicant(s): S. Kaxiras et al.
Case: 2-2
Serial No.: 09/591,918
Filing Date: June 9, 2000
Group: 2187
Examiner: Glenn A. Gossage

I hereby certify that this paper is being deposited on this date with the U.S. Postal Service as first class mail addressed to the Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.

Signature: V. Benvenuti Date: September 9, 2004

Title: Directory-Based Prediction Methods and Apparatus
for Shared-Memory Multiprocessor Systems

TRANSMITTAL OF APPEAL BRIEF

Mail Stop Appeal Brief - Patents
Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Sir:

Submitted herewith are the following documents relating to the above-identified patent application:

- (1) Appeal Brief in triplicate (original and two copies); and
- (2) Copy of Notice of Appeal, filed on July 6, 2004, with copy of stamped return postcard indicating receipt of Notice by PTO on July 9, 2004.

There is an additional fee of \$330 due in conjunction with this submission under 37 CFR §1.17(c). Please charge **Ryan, Mason & Lewis, LLP Account No. 50-0762** the amount of \$330, to cover this fee. In the event of non-payment or improper payment of a required fee, the Commissioner is authorized to charge or to credit **Deposit Account No. 50-0762** as required to correct the error. A duplicate copy of this letter and two copies of the Appeal Brief are enclosed.

Respectfully submitted,

Joseph B. Ryan

Date: September 9, 2004

Joseph B. Ryan
Attorney for Applicant(s)
Reg. No. 37,922
Ryan, Mason & Lewis, LLP
90 Forest Avenue
Locust Valley, NY 11560
(516) 759-7517



Kaxiras 2-2

THE UNITED STATES PATENT AND TRADEMARK OFFICE

Patent Application

Applicant(s): S. Kaxiras et al.
Case: 2-2
Serial No.: 09/591,918
Filing Date: June 9, 2000
Group: 2187
Examiner: Glenn A. Gossage

I hereby certify that this paper is being deposited on this date with the U.S. Postal Service as first class mail addressed to the Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.

Signature: [Signature] Date: September 9, 2004

Title: Directory-Based Prediction Methods and Apparatus
for Shared-Memory Multiprocessor Systems

APPEAL BRIEF

Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Sir:

Applicants hereby appeal the final rejection dated April 6, 2004 of claims 1-34 of the above-identified application.

REAL PARTY IN INTEREST

The present application is assigned to Agere Systems Inc. The assignee Agere Systems Inc. is the real party in interest.

RELATED APPEALS AND INTERFERENCES

There are no known related appeals or interferences.

09/13/2004 CCHAU1 00000077 500762 09591918

01 FC:1402 330.00 DA

STATUS OF CLAIMS

The present application was filed on June 9, 2000 with claims 1-34. Claims 1-34 are currently pending in the application. Claims 1, 22 and 28 are the independent claims.

Claims 1-4, 10, 21-25, 28-31 and 34 stand rejected under 35 U.S.C. §102(b). Claims 5-9, 11-20, 26, 27, 32 and 33 stand rejected under 35 U.S.C. §103(a). Claims 1-34 are appealed.

STATUS OF AMENDMENTS

There have been no amendments filed subsequent to the final rejection.

SUMMARY OF INVENTION

The present invention provides techniques for determining a set of predicted readers of a data block in a multiprocessor system. The invention generally involves determining a current set of readers of a data block which is subject to a write request, and generating the set of predicted readers based on the current set of readers and at least one additional set of readers representative of at least a portion of a global history of a directory associated with the data block.

An illustrative embodiment of the invention is in the form of a multiprocessor system 100 comprising nodes A, B and C coupled to an interconnection network 102, as shown in FIGS. 1 and 2 of the drawings. One or more of these nodes, or their respective processors 106A, 106B and 106C, at a particular point in time may be operating as a reader or a writer of a given data block. The nodes A, B and C include respective memories 108A, 108B and 108C, each of which includes an associated directory. See the specification at page 5, line 3, to page 6, line 8.

The claimed arrangements utilize a global history of a directory in generating a set of predicted readers of a data block which is subject to a write request. This aspect of the invention is described in the following manner at page 6, lines 18-26, of the specification:

The present invention in an illustrative embodiment provides a directory-based prediction mechanism which predicts the next set of readers of a block when a write request goes from the writer to the directory associated with the block. The mechanism predicts a likely set of readers of the value produced by the writer, and after the writer has finished

writing, this prediction is used to forward the data to all the predicted readers. Unlike conventional predictors which distinguish among blocks or among instructions to keep separate histories for blocks in the system, the prediction mechanism of the present invention merges together multiple sets of readers for multiple blocks served by the directory. This information is referred to herein as the global history of the directory.

Additional description relating to the global history of a directory is provided at, for example, page 10, lines 19-25, of the specification.

The present invention provides a number of significant advantages over conventional techniques. For example, the prediction techniques of the illustrative embodiment, through use of the global history of a directory, “provide improved prediction accuracy, both in terms of fewer false positives and fewer false negatives, relative to conventional techniques.” See the specification at page 4, lines 5-7.

ISSUES PRESENTED FOR REVIEW

1. Whether claims 1-4, 10, 21-25, 28-31 and 34 are anticipated under 35 U.S.C. §102(b) by S. Kaxiras, “Identification and Optimization of Sharing Patterns for Scalable Shared-Memory Multiprocessors” Ph.D. Thesis, Computer Sciences, University of Wisconsin-Madison, 1998 (hereinafter “Kaxiras”), which was cited by Applicants on their Information Disclosure Statement filed concurrently with the present application. Applicants note that the author S. Kaxiras of this primary reference and inventor Stefanos Kaxiras named on the present application are one and the same person.

2. Whether claims 5-7, 11-20, 26, 27, 32 and 33 are unpatentable under 35 U.S.C. §103(a) over Kaxiras.

3. Whether claims 8 and 9 are unpatentable under 35 U.S.C. §103(a) over Kaxiras in view of U.S. Patent No. 6,032,228 (hereinafter “Islam”).

GROUPING OF CLAIMS

With regard to Issue 1, claims 1-4, 10, 21-25, 28-31 and 34 stand or fall together.

With regard to Issue 2, claims 5, 26 and 32 stand or fall together, claims 6, 27 and 33 stand or fall together, claim 7 stands or falls alone, claims 11, 12 and 13 stand or fall together, claim 14 stands or falls alone, claim 15 stands or falls alone, claim 16 stands or falls alone, claim 17 stands or falls alone, claim 18 stands or falls alone, claim 19 stands or falls alone, and claim 20 stands or falls alone.

With regard to Issue 3, claim 8 stands or falls alone, and claim 9 stands or falls alone.

ARGUMENT

Issue 1

The Manual of Patent Examining Procedure (MPEP), Eight Edition, August 2001, §2131, specifies that a given claim is anticipated “only if each and every element as set forth in the claim is found, either expressly or inherently described, in a single prior art reference,” citing Verdegaal Bros. v. Union Oil Co. of California, 814 F.2d 628, 631, 2 USPQ2d 1051, 1053 (Fed. Cir. 1987). Moreover, MPEP §2131 indicates that the cited reference must show the “identical invention . . . in as complete detail as is contained in the . . . claim,” citing Richardson v. Suzuki Motor Co., 868 F.2d 1226, 1236, 9 USPQ2d 1913, 1920 (Fed. Cir. 1989).

The present invention as set forth in claim 1 is directed to a method of determining a set of predicted readers of a data block in a multiprocessor system. The method includes the steps of determining a current set of readers of a data block which is subject to a write request, and generating the set of predicted readers based on the current set of readers and at least one additional set of readers representative of at least a portion of a global history of a directory associated with the data block.

It is particularly important to note that the claimed arrangements utilize a global history of a directory. As indicated previously, this aspect of the invention is described in the following manner at page 6, lines 18-26, of the specification, with emphasis supplied:

The present invention in an illustrative embodiment provides a directory-based prediction mechanism which predicts the next set of readers of a block when a write request goes from the writer to the directory associated with the block. The mechanism predicts a

likely set of readers of the value produced by the writer, and after the writer has finished writing, this prediction is used to forward the data to all the predicted readers. Unlike conventional predictors which distinguish among blocks or among instructions to keep separate histories for blocks in the system, the prediction mechanism of the present invention merges together multiple sets of readers for multiple blocks served by the directory. This information is referred to herein as the global history of the directory.

Again, additional description relating to the global history of a directory is provided at, for example, page 10, lines 19-25, of the specification.

Applicants respectfully submit that Kaxiras fails to teach or suggest the claimed prediction arrangements involving use of a global history of a directory associated with a data block, and thus fails to provide the associated advantages of the type described at page 4, lines 5-7 of the specification.

The Examiner in formulating the §102(b) rejection of claim 1 over Kaxiras relies on the discussion at pages 206-207 thereof, and more particularly relies on the intersection-prediction approach described at page 207, section 2. This intersection-prediction approach is described in Kaxiras as follows:

The predictor is updated when the producer invalidates a sharing list and the identities of the consumers are collected on a temporary bit-map. The logical AND of the temporary bit-map and the predictor entry bit-map (that contains the consumers of the previous store-miss or store-write-fault) constitutes the prediction bit-map. After the prediction bit-map is calculated, the temporary bit-map is installed over the predictor entry's bit-map.

There is no mention whatsoever in this description regarding the claimed utilization of a global history of a directory associated with a data block. In fact, it appears that the cited passage is more properly viewed as simply an example of one of the “conventional predictors” referred to in the previously-quoted description from page 6, lines 18-26, of the specification. The other portions of

Kaxiras relied upon by the Examiner similarly fail to teach or suggest the claimed utilization of a global history of a directory.

In the final Office Action, at page 9, second paragraph, the Examiner argues that the “at least a portion of” language which precedes the term “global history” in claim 1 renders Kaxiras anticipatory, because “information regarding the sets of different consumers may be considered to represent a ‘global history.’” Applicants respectfully disagree. The argument proffered by the Examiner ignores the fact that the specification defines the term “global history” in a manner which distinguishes conventional predictors such as that described in Kaxiras. For example, as indicated above, the specification at page 6, lines 18-26, includes the following recitation:

Unlike conventional predictors which distinguish among blocks or among instructions to keep separate histories for blocks in the system, the prediction mechanism of the present invention merges together multiple sets of readers for multiple blocks served by the directory. This information is referred to herein as the global history of the directory.

The Examiner is in effect taking a conventional predictor and alleging that it anticipates “at least a portion of a global history of a directory” as claimed. However, since there is no global history of a directory described in Kaxiras, Kaxiras cannot be reasonably construed as being anticipatory of a portion of such a global history.

Claim 1 therefore includes limitations which are not taught or suggested by Kaxiras. The anticipation rejection of claim 1 is thus believed to be improper and should be withdrawn.

Independent claims 22 and 28 are not anticipated by Kaxiras for reasons similar to those identified above with regard to claim 1.

Dependent claims 2-4, 10, 21, 23-25, 29, 30, 31 and 34 are believed allowable for at least the reasons identified above with regard to independent claim 1.

Issue 2

Claims 5-7, 11-20, 26, 27, 32 and 33 stand rejected under §103(a) as being unpatentable over Kaxiras. Applicants respectfully traverse. These dependent claims are allowable for at least the reasons identified above with regard to their respective independent claims, and are believed to specify additional separately-patentable subject matter over Kaxiras.

With regard to dependent claims 5, 26 and 32, each of these claims specifies that the global history of the directory comprises a plurality of sets of previous readers processed by the directory, with the total number of the plurality of sets of previous readers corresponding to a designated history depth associated with generation of the set of predicted readers. As outlined above, there is no mention whatsoever in Kaxiras regarding a global history of a directory, as that term is defined in the specification, much less the particular global history limitations recited in dependent claims 5, 26 and 32.

The Examiner acknowledges that Kaxiras fails to meet the particular limitations recited in claims 5, 26 and 32, but argues that such limitations would be obvious in view of Kaxiras. In support of his argument, the Examiner provides the following statement regarding obviousness, at pages 9-10, section 5, of the final Office Action, with emphasis supplied:

Kaxiras . . . does not teach that the history information comprise a plurality of sets of previous readers processed by the directory. However, one of ordinary skill in the art would readily recognize that maintaining more than one set of previous readers in the history information allows one to better track the history of the previous readers or consumers and improve the prediction ability of the prediction scheme. Accordingly, it would have been readily obvious to one of ordinary skill in the art at the time the claimed invention was made to store of [sic] a plurality of sets of previous readers in the history information of Kaxiras in order to improve the prediction ability and thereby improve system performance by more accurately predicting where to “pre-send” information.

Applicants submit that this is a conclusory statement of obviousness, and insufficient to support the proposed modification of the reference teachings.

The Federal Circuit has stated that when patentability turns on the question of obviousness, the obviousness determination “must be based on objective evidence of record” and that “this precedent has been reinforced in myriad decisions, and cannot be dispensed with.” In re Sang-Su Lee, 277 F.3d 1338, 1343 (Fed. Cir. 2002). Moreover, the Federal Circuit has stated that “conclusory statements” by an examiner fail to adequately address the factual question of motivation, which is material to patentability and cannot be resolved “on subjective belief and unknown authority.” Id. at 1343-1344. There has been no showing in the present §103(a) rejection of objective evidence of record that would motivate one skilled in the art to modify Kaxiras in the manner alleged by the Examiner. The above-quoted statement of obviousness given by the Examiner in the final Office Action is precisely the type of subjective, conclusory statement that the Federal Circuit has indicated provides insufficient support for an obviousness rejection. It appears, in view of the above-quoted statement of obviousness provided by the Examiner, that the Examiner in rejecting claims 5, 26 and 32 over Kaxiras has simply undertaken a hindsight-based piecemeal reconstruction of the claimed invention based on the disclosure provided by Applicants. Such an approach is improper.

With regard to dependent claims 6, 27 and 33, each of these claims specifies that the global history is maintained in a shift register having a number of storage locations corresponding to the designated history depth. Kaxiras fails to teach or suggest a global history of a directory, much less the claimed use of a shift register to implement such a history. Thus, Kaxiras does not render obvious these particular limitations.

With regard to dependent claim 7, this claim calls for a global history having a particular history depth, namely, a history depth greater than two. Kaxiras fails to teach or suggest a global history of a directory, much less the claimed use of a history depth greater than two. The Examiner in the final Office Action at page 10, last paragraph, to page 11, first paragraph, provides only a conclusory statement of obviousness, not based on any objective evidence of record. *A prima facie* case of obviousness of claim 7 over Kaxiras has therefore not been established.

With regard to dependent claims 11, 12 and 13, these claims are believed allowable for at least the reasons identified above with regard to independent claim 1.

With regard to dependent claim 14, this claim specifies that a dynamically-selected function, applied to a current set of readers and at least one additional set of readers, is selected on a per-

program basis, such that each of a plurality of programs running in a multiprocessor system can independently determine the function to be applied to determine the set of predicted readers. The Examiner at page 11, second paragraph, simply states that the limitations “would have been readily obvious” without providing any objective evidence in support of such an argument. Again, a *prima facie* case has not been established.

With regard to dependent claim 15, this claim specifies that a dynamically-selected function, applied to a current set of readers and at least one additional set of readers, is selected under program control and can be selected at run time by a given program running on a multiprocessor system. The Examiner again relies entirely on subjective reasoning, and fails to cite any particular portion of the Kaxiras reference which is alleged to suggest the limitation in question.

With regard to dependent claim 16, this claim specifies that a dynamically-selected function, applied to a current set of readers and at least one additional set of readers, is selected on a per-page basis, such that the function applied can be determined independently for each of a plurality of memory pages, each of which may be comprised of one or more data blocks. Applicants submit that there is no teaching or suggestion of such a limitation in Kaxiras. The Examiner has again simply relied on subjective reasoning, and fails to cite any particular portion of the Kaxiras reference which is alleged to suggest the limitation in question.

With regard to dependent claim 17, this claim specifies that a dynamically-selected function, applied to a current set of readers and at least one additional set of readers, is selected based at least in part on information regarding network utilization. Applicants have been unable to locate any suggestion of this limitation in Kaxiras, and the Examiner has not cited any particular portion of Kaxiras in support of his argument. It is therefore believed that a proper *prima facie* case has not been established by the Examiner.

With regard to dependent claim 18, this claim specifies that the generating step of claim 1 further includes utilizing information regarding the global history of the directory in conjunction with at least a subset of cache address information associated with one or more of the readers to determine the set of predicted readers. The Examiner relies on the “hybrid prediction scheme” of Kaxiras, but fails to provide anything but a subjective, conclusory statement regarding the claimed use of cache address information in conjunction with information regarding a global history of a directory. See

the final Office Action at page 11, last paragraph, to page 12, first paragraph. Applicants respectfully submit that the Examiner has failed to establish a proper *prima facie* case of obviousness of claim 18 over Kaxiras.

With regard to dependent claim 19, this claim specifies that the generating step of claim 1 further includes utilizing information regarding the global history of the directory in conjunction with processor node information associated with one or more of the readers to determine the set of predicted readers. Again, the Examiner relies on the “hybrid prediction scheme” of Kaxiras, but fails to provide anything but a subjective, conclusory statement regarding the claimed use of processor node information in conjunction with information regarding a global history of a directory. It is believed that the Examiner has failed to establish a proper *prima facie* case of obviousness of claim 19 over Kaxiras.

With regard to dependent claim 20, this claim specifies that the generating step of claim 1 further includes utilizing information regarding the global history of the directory in conjunction with program counter information associated with one or more of the readers to determine the set of predicted readers. The Examiner relies on the same subjective, conclusory statement relied upon in the context of claims 18 and 19, and fails to cite to any objective evidence of record indicating some suggestion regarding the claimed use of program counter information. Accordingly, a proper *prima facie* case of obviousness of claim 20 over Kaxiras has not been established.

Issue 3

Claims 8 and 9 stand rejected under §103(a) as being unpatentable over Kaxiras in view of Islam. Applicants respectfully traverse. The Islam reference fails to supplement the fundamental deficiencies of Kaxiras as applied to the independent claims, in that it fails to provide any teaching or suggestion regarding the claimed utilization of a global history of a directory.

With regard to the §103(a) rejection over Kaxiras and Islam, Applicants note that a proper *prima facie* case of obviousness over a pair of references requires that the cited references when combined must “teach or suggest all the claim limitations,” and that there be some suggestion or motivation, either in the references themselves or in the knowledge generally available to one of

ordinary skill in the art, to combine the references or to modify the reference teachings. See Manual of Patent Examining Procedure (MPEP), Eighth Edition, August 2001, §706.02(j).

Applicants submit that the Examiner has failed to establish a proper *prima facie* case of obviousness in the present §103(a) rejection, in that the Kaxiras and Islam references, even if assumed to be combinable, fail to teach or suggest all the claim limitations, and in that no cogent motivation has been identified for combining the references or modifying the reference teachings to reach the claimed invention.

Dependent claim 8 specifies that each reader in the system maintains an accessed bit for each of a plurality of data blocks, the accessed bit of a particular reader for a given data block indicating whether the particular reader has actually read the given data block.

The Examiner acknowledges that this limitation is not shown in Kaxiras, but argues that it is met by the proposed combination of Kaxiras and Islam. More particularly, the Examiner relies on the “accessed bit” disclosed in column 3, lines 44-47, of Islam. However, the relied-upon accessed bit of Islam is clearly described therein as being maintained for a particular cache entry. Applicants respectfully submit that an accessed bit for a cache entry is not an accessed bit of the type claimed, the latter explicitly described as being maintained by a reader for a given data block indicating whether that particular reader has read the data block. Instead, the accessed bit for the cache entry in Islam simply indicates whether or not that cache entry has been accessed by any reader. In other words, there is no reader-specific information maintained by the accessed bit disclosed in Islam. The combined teachings of Kaxiras and Islam therefore fail to meet the limitations of claim 8.

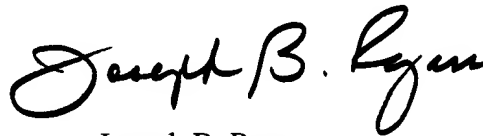
The Examiner states that it would be obvious to combine Kaxiras and Islam “in order to obtain feedback on the actual use of the data and thereby improve system performance.” Again, this is a conclusory statement of obviousness of the type prohibited by the In re Sang-Su Lee case cited above. In addition, the proposed combination would clearly fail to provide the benefits associated with the reader-specific information of the claimed accessed bit.

With regard to dependent claim 9, this claim specifies that accessed bit information is sent by the particular reader to the directory in conjunction with an invalidation acknowledgment. There is no teaching in the proposed combination regarding this particular limitation. The accessed bit of

Islam, relied upon by the Examiner, does not indicate whether a particular reader has read a corresponding data block, nor is it sent anywhere by any particular reader. All that the Examiner provides in formulating the §103(a) rejection of claim 9 is another conclusory statement of obviousness, to the effect that the limitations would be obvious “so as to reduce the amount [sic] coherence traffic.”

In view of the above, Applicants believe that claims 1-34 are in condition for allowance, and respectfully request the withdrawal of the §102(b) and §103(a) rejections.

Respectfully submitted,

A handwritten signature in black ink that reads "Joseph B. Ryan". The signature is written in a cursive, flowing style with a large, prominent "J" and "R".

Date: September 9, 2004

Joseph B. Ryan
Attorney for Applicant(s)
Reg. No. 37,922
Ryan, Mason & Lewis, LLP
90 Forest Avenue
Locust Valley, NY 11560
(516) 759-7517

APPENDIX

1. A method of determining a set of predicted readers of a data block in a multiprocessor system, the method comprising the steps of:

determining a current set of readers of a data block which is subject to a write request;

and

generating the set of predicted readers based on the current set of readers and at least one additional set of readers representative of at least a portion of a global history of a directory associated with the data block.

2. The method of claim 1 wherein the generating step further comprises the step of applying a function to the current set of readers and at least one additional set of readers.

3. The method of claim 2 wherein the function comprises at least one of a union function, an intersection function and a pattern-based function.

4. The method of claim 1 wherein the directory and the data block comprise elements of a memory associated with a processor node of the multiprocessor system.

5. The method of claim 1 wherein the global history of the directory comprises a plurality of sets of previous readers processed by the directory, the total number of the plurality of sets of previous readers corresponding to a designated history depth associated with generation of the set of predicted readers.

6. The method of claim 5 wherein the global history is maintained in a shift register having a number of storage locations corresponding to the designated history depth.

7. The method of claim 6 wherein the history depth is greater than two.

8. The method of claim 1 wherein each reader in the system maintains an accessed bit for each of a plurality of data blocks, the accessed bit of a particular reader for a given data block indicating whether the particular reader has actually read the given data block.

9. The method of claim 8 wherein accessed bit information is sent by the particular reader to the directory in conjunction with an invalidation acknowledgment.

10. The method of claim 1 wherein after the requested write on the data block is completed, the resulting data block is sent to each of the readers in the set of predicted readers.

11. The method of claim 1 wherein the generating step utilizes a direct update mechanism.

12. The method of claim 1 wherein the generating step utilizes a forwarded update mechanism.

13. The method of claim 2 wherein the function is selected dynamically.

14. The method of claim 13 wherein the function is selected on a per-program basis, such that each of a plurality of programs running in the multiprocessor system can independently determine the function to be applied to determine the set of predicted readers.

15. The method of claim 13 wherein the function is selected under program control and can be selected at run time by a given program running on the multiprocessor system.

16. The method of claim 13 wherein the function is selected on a per-page basis, such that the function applied can be determined independently for each of a plurality of memory pages, each of which may be comprised of one or more data blocks.

17. The method of claim 13 wherein the function is selected based at least in part on information regarding network utilization.

18. The method of claim 1 wherein the generating step further includes utilizing information regarding the global history of the directory in conjunction with at least a subset of cache address information associated with one or more of the readers to determine the set of predicted readers.

19. The method of claim 1 wherein the generating step further includes utilizing information regarding the global history of the directory in conjunction with processor node information associated with one or more of the readers to determine the set of predicted readers.

20. The method of claim 1 wherein the generating step further includes utilizing information regarding the global history of the directory in conjunction with program counter information associated with one or more of the readers to determine the set of predicted readers.

21. The method of claim 1 wherein each of at least a subset of the readers corresponds to a particular processor node in the multiprocessor system.

22. An apparatus for determining a set of predicted readers of a data block in a multiprocessor system, the apparatus comprising:

a processor node operative to determine a current set of readers of a data block which is subject to a write request, and to implement a prediction mechanism which generates the set of predicted readers based on the current set of readers and at least one additional set of readers representative of at least a portion of a global history of a directory associated with the data block.

23. The apparatus of claim 22 wherein the set of predicted readers is generated at least in part by applying a function to the current set of readers and at least one additional set of readers.

24. The apparatus of claim 23 wherein the function comprises at least one of a union function, an intersection function and a pattern-based function.

25. The apparatus of claim 22 wherein the directory and the data block comprise elements of a memory associated with the processor node.

26. The apparatus of claim 22 wherein the global history of the directory comprises a plurality of sets of previous readers processed by the directory, the total number of the plurality of sets of previous readers corresponding to a designated history depth associated with generation of the set of predicted readers.

27. The apparatus of claim 26 wherein the global history is maintained in a shift register having a number of storage locations corresponding to the designated history depth.

28. A multiprocessor system comprising:
a plurality of processor nodes, at least a given one of the processor nodes being operative to determine a current set of readers of a data block which is subject to a write request, the given processor node implementing a prediction mechanism which generates a set of predicted readers of the data block based on the current set of readers and at least one additional set of readers representative of at least a portion of a global history of a directory associated with the data block.

29. The system of claim 28 wherein the set of predicted readers are generated at least in part by applying a function to the current set of readers and at least one additional set of readers.

30. The system of claim 29 wherein the function comprises at least one of a union function, an intersection function and a pattern-based function.

31. The system of claim 28 wherein the directory and the data block comprise elements of a memory associated with the given processor node.

32. The system of claim 28 wherein the global history of the directory comprises a plurality of sets of previous readers processed by the directory, the total number of the plurality of sets of previous readers corresponding to a designated history depth associated with generation of the set of predicted readers.

33. The system of claim 32 wherein the global history of the directory is maintained in a shift register associated with a corresponding one of the processor nodes and having a number of storage locations corresponding to the designated history depth.

34. The system of claim 28 wherein each of at least a subset of the readers corresponds to a particular one of the processor nodes in the multiprocessor system.